

IN THE CLAIMS:

Please cancel claims 32-33. Please add new claims 34-37 and amend claims 1, 2, 6, 7, 11, 12, 16, 17, 21, 22, 26, 27 and 31 as shown below.

- AI
1. (Currently Amended) A data storage system comprising:

a first volume having a first storage volume characteristic;

a second volume having a second storage volume characteristic; and

a computing node coupled to said first volume and said second volume, wherein said computing node includes a file system for identifying files stored by said first volume and said second volume;

wherein said file system includes a directory structure having a directory allocated to said first volume and said second volume, wherein said directory includes having an entry corresponding to a file maintained by said file system, and wherein said entry includes a field containing a volume identifier indicative of which of said first or said second volumes said file is stored within.
 2. (Currently Amended) The system as recited in claim 1, wherein said file system is configured to allocate space on said first volume in response to receiving a request specifying said first storage volume characteristic and said second volume in response to receiving a request specifying a said second storage volume characteristic from a software application.
 3. (Original) The system as recited in claim 2, wherein each of said first volume and said second volume comprises a single storage device.

AI

4. (Original) The system as recited in claim 2, wherein each of said first volume and said second volume comprises a multiple storage device system.

5. (Original) The system as recited in claim 4, wherein said multiple storage device system is a redundant array of inexpensive disks (RAID) storage system.

6. (Currently Amended) A file system for use in a computing node coupled to a first volume and a second volume, wherein said file system is configured to identify files stored by said first volume and said second volume, wherein said file system includes a directory structure having a directory allocated to said first volume and said second volume, wherein said directory includes having an entry corresponding to a file maintained by said file system, and wherein said entry includes a field containing a volume identifier indicative of which of said first or said second volumes said file is stored within.

7. (Currently Amended) The file system as recited in claim 6, wherein said file system is further configured to allocate space on said first volume in response to receiving a request specifying a first storage volume characteristic and said second volume in response to receiving a request specifying a second storage volume characteristic from a software application.

8. (Original) The file system as recited in claim 7, wherein each of said first volume and said second volume comprises a single storage device.

9. (Original) The file system as recited in claim 8, wherein each of said first volume and said second volume comprises a multiple storage device system.

10. (Original) The file system as recited in claim 9, wherein said multiple storage device system is a redundant array of inexpensive disks (RAID) storage system.

AI
11. (Currently Amended) A method of operating a file system which identifies files stored by a first volume and a second volume, said method comprising:

providing a filename corresponding to a file maintained by said file system; and

accessing an entry in a directory ~~structure~~ allocated to said first volume and said second volume, wherein said entry includes a field containing a volume identifier indicative of which of said first or said second volumes said file is stored within.

12. (Currently Amended) The method as recited in claim 11, wherein said method further comprises allocating space on said first volume in response to receiving a request specifying a first storage volume characteristic and said second volume in response to receiving a request specifying a second storage volume characteristic from a software application.

13. (Original) The method as recited in claim 12, wherein each of said first volume and said second volume ~~comprises~~ a single storage device.

14. (Original) The method as recited in claim 12, wherein said first volume and said second volume are each a logical volume, wherein said each logical volume comprises a multiple storage device system.

15. (Original) The method as recited in claim 14, wherein said multiple storage device system is a redundant array of inexpensive disks (RAID) storage system.

16. (Currently amended) A computer readable medium comprising instructions for operating a file system which identifies files stored by a first volume and a second volume, wherein said instructions are executable by a computing node to implement a method comprising:

providing a filename corresponding to a file maintained by said file system; and

AI
accessing an entry in a directory ~~structure~~ allocated to said first volume and said second volume, wherein said entry includes a field containing a volume identifier indicative of which of said first or said second volumes said file is stored within.

17. (Currently Amended) The computer readable medium as recited in claim 16, wherein said method further comprises allocating space on said first volume in response to receiving a request specifying a first storage volume characteristic and said second volume in response to receiving a request specifying a second storage volume characteristic from a software application.

18. (Original) The computer readable medium as recited in claim 17, wherein each of said first volume and said second volume comprises a single storage device.

19. (Original) The computer readable medium as recited in claim 17, wherein each of said first volume and said second volume comprises a multiple storage device system.

20. (Original) The computer readable medium as recited in claim 19, wherein said multiple storage device system is a redundant array of inexpensive disks (RAID) storage system.

SUB B17 21. (Currently Amended) A data storage system comprising:

a first volume having a first storage volume characteristic;

a second volume having a second storage volume characteristic; and

a computing node coupled to said first volume and said second volume, wherein said computing node includes a file system for identifying a first file

AI
stored on said first volume and a second file stored on said second volume;

wherein said file system includes a directory structure having a directory which includes a first entry corresponding to said first file and a second entry corresponding to said second file.

22. (Currently Amended) The system as recited in claim 21, wherein said file system is configured to allocate space on said first volume in response to receiving a request specifying said first storage volume characteristic and said second volume in response to receiving a request specifying a said second storage volume characteristic from a software application.

23. (Original) The system as recited in claim 22, wherein each of said first volume and said second volume comprises a single storage device.

24. (Original) The system as recited in claim 22, wherein each of said first volume and said second volume comprises a multiple storage device system.

25. (Original) The system as recited in claim 24, wherein said multiple storage device system is a redundant array of inexpensive disks (RAID) storage system.

26. (Currently Amended) A method comprising:

storing a first file on a first volume having a first storage volume characteristic based on a first set of storage characteristics desired for said first file, wherein said first file is located in a directory of a directory structure maintained by a file system; and

storing a second file on a second volume having a second storage volume characteristic based on a second set of storage characteristics desired for said second file, wherein said ~~first~~ second file is located in said directory.

A1 27. (Currently Amended) The method as recited in claim 26, wherein said method further comprises allocating space on said first volume in response to receiving a request specifying said storage volume characteristic and said second volume in response to receiving a request specifying a said second storage volume characteristic from a software application.

28. (Original) The method as recited in claim 27, wherein each of said first volume and said second volume comprises a single storage device.

29. (Original) The method as recited in claim 27, wherein said first volume and said second volume are each a logical volume, wherein said each logical volume comprises a multiple storage device system.

30. (Original) The method as recited in claim 29, wherein said multiple storage device system is a redundant array of inexpensive disks (RAID) storage system.

31. (Currently Amended) A computer memory containing a directory structure maintained by a file system having a first entry in a directory corresponding to a first file and a second entry in said directory corresponding to a second file, wherein said first file is stored on a first volume having a first set of storage characteristics and said second file is stored on a second volume having a second set of storage characteristics.

32. (Cancelled)

33. (Cancelled)

SUB B37

34. (New) The system as recited in claim 1, wherein said entry includes another field containing an index number associated with metadata corresponding to said file.

A2

35. (New) The system as recited in claim 34, wherein said first volume and said second volume each specify a set of methods for manipulating said metadata and for allocating data blocks.

36. (New) The file system as recited in claim 6, wherein said entry includes another field containing an index number associated with metadata corresponding to said file.

37. (New) The file system as recited in claim 36, wherein said first volume and said second volume each specify a set of methods for manipulating said metadata and for allocating data blocks.